Mathematical Statistics
Stockholm University

A classical dataset from Williams, and its role in the study of supersaturated designs.

Rolf Sundberg

Research Report 2007:22

Postal address:
Mathematical Statistics
Dept. of Mathematics
Stockholm University
SE-106 91 Stockholm
Sweden

Internet:
http://www.math.su.se/matstat

A classical dataset from Williams, and its role in the study of supersaturated designs.

Rolf Sundberg[*]

December 2007

Abstract

A Plackett–Burman type dataset from a paper by Williams (1968), with 28 observations and 24 two-level factors, has become a standard dataset for illustrating construction (by halving) of supersaturated designs (SSDs) and for a corresponding data analysis. The aim here is to point out that for several reasons this is an unfortunate situation. The original paper by Williams contains several errors and misprints. Some are in the design matrix, which will here be reconstructed, but worse is an outlier in the response values, which can be observed when data are plotted against the dominating factor. In addition, the data should better be analysed on log-scale than on original scale. The implications of the outlier for SSD analysis are drastic, and it will be concluded that the data should be used for this purpose only if the outlier is properly treated (omitted or modified).

Key words: half-fraction, log-transformation, outlier, Plackett–Burman, SSD

[*]Postal address: Mathematical Statistics, Stockholm University, SE-106 91, Sweden. E-mail: rolfs@math.su.se. Website: www.math.su.se/∼rolfs

1 Introduction

A classical dataset, from an article by Williams in 1968 [1], will be revisited. This dataset, of Plackett–Burman type with $n = 28$ observations and $k = 24$ factors has been used as an exercise by Box & Draper [2, 3] and for trying variable selection strategies [4]. Of more importance here, however, is that it has become very popular for trying construction and analysis methods for supersaturated designs (SSD), ever since a half-fraction of it was used by Lin [5]. Quoting [6]: *"It has, in a sense, become the default data set by which to judge SSD analysis methods"*. Supersaturated designs are characterized by having $n \leq k$, and can be useful when only a few factors are believed to be of importance (the effects sparsity principle), in particular in ruggedness tests when none of the factors is expected to be very influential.

We will take a closer look at this dataset. It suffers from many typing errors in original form [1], and we will reconstruct its complete design matrix. We will argue that data should be regarded on log-scale, but also demonstrate the presence of an outlier, whose influence on the SSDs is serious. This has led to several misinterpretations of the properties of proposed methods for statistical analysis of SSD data. These comments were specifically triggered by a recent paper [7] where this dataset is used to investigate the properties of the so-called FEAR method of SSD data analysis.

2 History

The dataset in question, in [7] misleadingly called Data set of Abraham [8], originates from a paper by Williams in 1968 [1]. According to this origin it was designed as a Plackett–Burman type experiment with 28 runs but restricted to 24 two-level factors. Two decades later it was used as Exercise 5.20 in the book by Box & Draper [2]. They identified and corrected a typographical sign error in the design matrix, but did not notice that two factors in the design matrix had identical sign columns. Lin [5] took up the dataset in order to try his method for generating SSDs, constructing a 14 runs half-fraction from the original 28 runs. Wang [9] pointed out that two sign columns in [1] (and [5]) were identical. In later studies one of these column duplicates has typically been deleted, with later factors renumbered. This has occasionally misled others to quote wrong factor numbers [10]. In [2], the neglect in combination with the primitive computation facilities at the time misled the authors to give an erroneous ANOVA table in the

solutions to the exercise. This particular error has been corrected in the 2nd edition [3], and is of no further concern here.

Since Lin's paper appeared in 1993 [5], Williams' dataset has been used in numerous papers on supersaturated designs. Some of these papers are of particular interest here. Lin ran a stepwise regression to select factors to include in the model. His method appeared successful, because he got essentially the same effects as Williams had obtained for the whole dataset. Wang [9] tried effect estimation in the same way as Lin, but on the complementary half-fraction of the data, and was surprised to find a quite different set of presumably significant effects from his stepwise regression. Abraham *et al.* [8] went a step further and tried four pairs of half-size SSDs. They did not refer to [9], but found similarly that different half-fractions yielded different choices of factors, They made two suggestions, "The nonorthogonality of the columns of the design matrix X is the root of the problem" (repeated in [11]), and "The dramatically different conclusions found in the eight different subsets appear to be mainly due to random variation in the data". Dejaegher *et al.* [7] refer to different confounding patterns and "a relatively high noise level". Selection error control has been studied in [12] and [13]. The former suggests that the disagreement between Lin's and Wang's analyses "could easily be attributable to the multiplicity problem", whereas the latter suggested that the reason might be that "Williams' data may not follow the effects sparsity principle well". Lu & Wu [10] mention that different half-fractions yield very different results but apparently avoid this problem by trying their own factor selection method only on Lin's half-fraction.

Hence, there is a large number of proposals concerning the reason for the problems with this data set. However, in none of these papers is a plot of the data to be found. It will here be demonstrated that Williams' dataset should be considered on the log-scale instead of original scale and that it suffers from an influential outlier, which is seriously influential for the small datasets constructed as SSDs. The latter consequences motivate the advice that these data should not be used to study SSD properties, except for pointing out this vulnerability—otherwise the outlier must be taken care of (omitted or modified).

3 Data structure

3.1 Design

Williams [1] denotes his design as "Plackett–Burman in 28 experiments", referring to the standard paper by Plackett & Burman [14]. Dejaegher *et al.* [7] raise a question as to this property, by categorically stating (p. 307) that it is not P–B. We will here first indicate how it can be true that Williams actually did use a Plackett–Burman design, and hence that the questioning in [7] is not motivated. At the same time we will (and have to) reconstruct columns 13 and 16, which by mistake were printed identical in Williams' paper, and the 3 supplementary contrast columns (mixed interactions) needed to have a set of 27 orthogonal effect estimates representing (together with the grand mean) uniquely the whole original dataset. A complete and revised sign table is given in the Appendix.

A standard Plackett–Burman construction of a $n = 28$ observations design matrix uses permutations of three 9×9 matrices for 27 rows supplemented by a row of only minus signs. Clearly the table in [2] does not have a row of minus signs. The first clue to understanding Williams is that for quantitative factors Box & Draper conventionally let the higher value correspond to the $+$ level. Williams [1], however, did not use the signs in this way. There are some inconsistencies in the levels themselves in [1], and in the ordering of the factors, but if we simply change level notations for six factors in the table from [2] so that the last row contains only minus signs, it turns out that we almost see the standard construction. The differences are only as follows:

- The 27×27 part of the full design matrix is (essentially—see below) the transpose of the matrix devised by Plackett and Burman, but one is obtained from the other by reversing the order of both rows and columns, and that makes an insignificant difference. If Williams were fully aware of this, or just lucky, is unclear from his paper.

- The sign column for factor 13 does not match the construction, whereas for factor 16 it does. A design-matching column 13 and the given column 16 actually yield the effects calculated by Williams, to a typical degree of precision (which is not very high in [1]). This strongly indicates that column 13 was misprinted in [1], whereas column 16 is right. The version given in the Appendix is corrected in this way.

4

- Design-matching column 24 appears in the expected position of column 21, and the design-matching columns 21–23 are moved one step to the right. Of course this does not matter for the character of the design.

- The last 3 design-matching columns, A1–A3 were not used by Williams. They correspond to mixed interaction effects.

The conclusion is that Williams' design is of P–B type. The interaction columns are important for the statistical inference. Lin [5] used the last column (A3) as defining contrast for his half-fraction construction. Abraham *et al.* [8] tried all three (A1–A3), and the revised column 13, in this way obtaining four pairs of different half-samples.

3.2   Data

Already a brief look at the column of response values y in Table 1 shows that the variation in magnitude is large, a factor 10 from the smallest ($y = 32$) to the largest ($y = 300$), and the distribution is clearly upwards skew. Whether the variation is random or systematic (due to some of the design factors), it is anyway natural to try modelling by multiplicative effects, rather than additive. Figures 1 a) and b) show data on original scale and on logarithmic scale, respectively, in combination with a split according to the clearly dominating factor 15. Surprisingly, there is no plot of these data to be found in the extensive previous literature using or discussing them. Factor 15 being dominant is not only seen from Figure 1, but also from ANOVA or regression analyses, where it has a much higher effect/coefficient than any other factor. In Figure 1 we see that after logarithmic transformation, the two subsets show variation of more similar size, except for one extreme observation when factor 15 is on its plus-level ($x_{15} = +1$), run #14. It appears obvious from Figure 1 that #14 is an outlier, which if not neutralized might have detrimental effects. For example, in the log-scale two-sample model represented by Figure 1b, but with run #14 excluded, its actual value deviates as much as $4.5\,\hat{\sigma}$ from the corresponding sample mean. Hence, we make two observations:

- Log-scale is more appropriate than the original scale.

- Run #14 is an outlier.

Further support is provided by the half-normal quantile plots in Figure 2, where estimated effects (the effect of factor 15 excluded) fit a half-normal

5
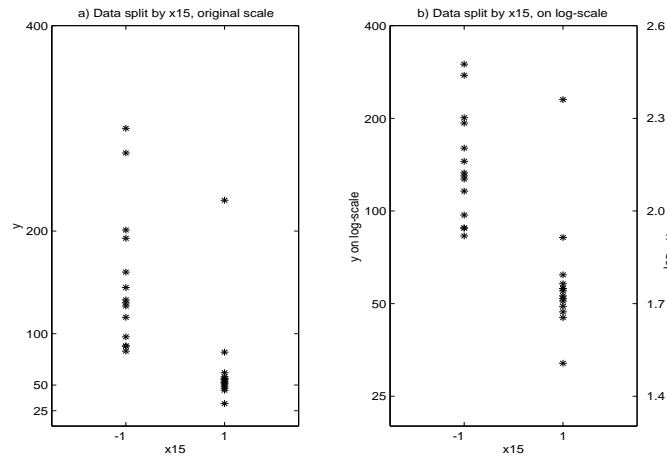
Figure 1: Scatter-plots of all data, split by the level of the dominating factor 15. The high value when $x_{15} = +1$ belongs to run #14.
a) Data on original scale.
b) Data on log-scale.
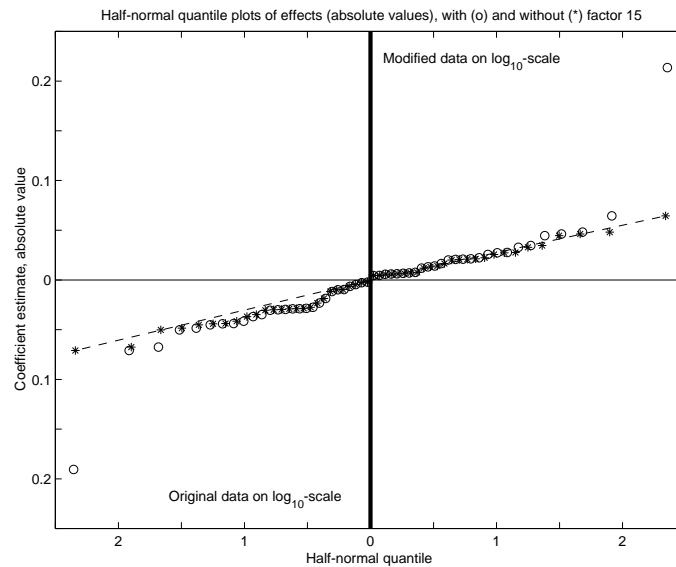
Figure 2: Four half-normal quantile plots of effect estimates.
Left part: Original data log-transformed.
Right part: Value for run #14 changed to the sample mean under $x_{15} = +1$.
Circles: All 27 effects included.
Asterisks: Effect of the influential factor 15 not included.
It is seen that the modified data fit the halfnormal better, and the asterisks then fit perfectly. This supports the conclusions that run #14 is an outlier and that factor 15 is the only certainly influential factor.

6

almost perfectly (right part), when data are on log-scale and run #14 has been replaced by its x_{15}-sample mean. The corresponding fit is not that good with the original value of run #14 (left part) and even worse with data on original scale. The latter is not shown here, but can be found in [7], except for the defect that the effects corresponding to the (main-effect-free) interaction columns (A1–A3 in the Appendix) have been omitted or forgotten.

The log transformation makes statistical conclusions based on the normal distribution (p-values and confidence intervals) more reliable. The more constant variance improves power. If several factors have effects, the transformation will also imply that their influence is differently modelled, multiplicatively instead of additively. In the present case, however, it cannot be concluded from data that any factor other than number 15 has an effect. Figure 2, for example, does not indicate that any other factor is influential, and a full main effects model yields factor 15 alone as significant. However, with all 24 factors in a main effects model, there are only 3 degrees of freedom, which means that ordinary t-tests for factor effects in this model are not very powerful. To be more liberal, we give in Table 1 below the first factors appearing in a forward selection. When judging p-values, there is also a need to compensate for multiple testing by demanding lower values than in a single test. For Lin's SSD data, this was stressed in [12], and the result was that only factor 15 was found significant.

On original scale, and with full data, the coefficient value for factor 15 is twice as high as the next largest coefficient. After going over to logarithms, and excluding the outlier, twice has increased to a factor 3.

| Orig. scale, full data | | | $\log_{10}$-scale, full data | | | $\log_{10}$-scale, obs. 14 excl. | | |
Factor	Coeff.	(p-value)	Factor	Coeff.	(p-value)	Factor	Coeff.	(p-value)
15	−43	(0.001)	**15**	−0.19	(0.000)	**15**	−0.21	(0.000)
17	−21	(0.03)	**17**	−0.07	(0.04)	**8**	+0.07	(0.01)
20	−24	(0.03)	**20**	−0.07	(0.04)	**24**	−0.05	(0.04)
						17	−0.05	(0.04)
						20	−0.04	(0.04)

Table 1. Forward selection main effects, as regression coefficients for x-variables at values $x = \pm 1$. Listed p-values are F-test inclusion probabilities, after rounding to two or three decimal digits. Coefficient estimates in the 3rd column are from the main-effects model with the five factors listed.

4 Consequences for SSD analysis studies based on Williams' data.

As mentioned above, Williams' data [1] have been used a lot, starting with Lin 1993 [5], for investigation and comparison of statistical analysis methods applied to supersaturated designs, mostly various forms of variable selection methods. Soon afterwards, Wang [9] tried stepwise regression on the complementary half-fraction and got such a different sequence of factors that he concluded that one of the two analyses must be misleading. He also suggested that from the given data alone, without using subject matter expertise, it would be impossible to tell which of them. His statement is not true, because the understanding of the full data given in the previous section, pointing out the extreme run #14, will explain his observation, in particular why factor 15 dominates in one half-fraction but does not even enter in the analysis of the complementary half. At the same time it explains analogous observations reported in [7, 13], but more generally it clarifies the wider, "not entirely encouraging" observation made by Abraham *et al.* [8], when they try four pairs of half-fractions and have problems understanding why factor 15 is entering or not entering in their variable selection procedure.

$\log_{10}$-scale, Wang half			$\log_{10}$-scale, Wang, run #14 excluded		
Factor	Coeff.	(p-value)	Factor	Coeff.	(p-value)
4	+0.16	(0.02)	**15**	−0.18	(0.000)
			24	−0.09	(0.03)

Table 2. Forward selection main effects, as regression coefficients for x-variables at values $x = \pm 1$. Listed p-values are F-test inclusion probabilities, and next p-value was > 0.05. Coefficient estimates are from the model with the factors given in the table.

Table 2 is intended to show that run #14 is the problem. It uses data on log-scale, but this is not crucial. We can see that Wang's half does not at all suggest factor 15, but as soon as run #14 is excluded from Wang's half-fraction, factor 15 will be highly significant again. To understand why factor 15 has lost its significance, it helps to contemplate how a t-test statistic for the main effect of factor 15 is affected when the full dataset is reduced to Wang's half-fraction. Its numerator will decrease substantially (by 26%), because the outlying observation 14 will have much higher influence with 7 than with 14 observations per sample. At the same time the pooled sample

standard deviation in the denominator will increase, for the same reason (by 21%). Finally, the formula for the standard error will increase by a factor $\sqrt{2}$ when sample sizes are reduced from 14 to 7. All these effects together reduce the t-test statistic by more than a factor 2, at the same time as the reduced degrees of freedom (from 26 to 12) would have required a higher t-value to keep the p-value level. The result is that factor 15 is not even 10% significant with original response and just below 5% significant on log-scale. Instead the potential to (falsely) explain run #14 by other factors will increase as samples sizes are reduced.

Lin's half-fraction happened to exclude run #14. The effect of this is similar to deliberate exclusion of the suspect run #14 from the full data, and the t-test statistic for factor 15 is substantially increased.

Similar effects will be seen whatever sign column is used as defining contrast for the half-fraction. If run #14 is in the fraction selected, the risk is high that factor 15 will not be significant even when tested alone, and that it will not enter in a selection procedure, whereas the opposite effect will be seen in the complementary half-fraction. This adequately explains the observations made in [8], when they try different branching columns (or defining contrasts).

If we believe that run #14 should be excluded from the complete dataset in order to achieve representative data, we can draw the following conclusions for the study of SSDs using Williams' data with run #14. Lin's and other half-fractions excluding run #14 show better results than the full set of all runs, whereas Wang's and corresponding half-fractions show misleadingly bad results. Earlier papers have not clarified this, but only referred to "random variation in the data" or other not very elucidative arguments (see Section 2 above).

5   Concluding discussion

Many old datasets are repeatably used for statistical method comparison and illustration. Some such datasets are frequently used for a period of time until it is pointed out that they are not suitable for the intended purpose, due to previously unrecognized features, for example lurking variables, design structure elements, outliers, etc. [15, 16]. Here it has been argued that due to a previously unnoticed outlier, Williams' data from 1968 [1] are likely to be misleading when trying methods of statistical analysis on subsets of

9

them.

The literature about outlier detection and identification in factorial experiments [17–21] was of limited use in this case. It is based on the redundancy of higher-order effect estimates, which is not equivalent with the sparsity of effects in a main effects model. In particular, we had to confine ourselves to the blunter half-normal plots instead of the usual (Daniels) normal quantile plot of effects.

The reason for the outlier cannot be reconstructed. It might be a misprint, a gross error, the wrong design point (not the intended level of factor 15), or just an otherwise extreme outcome. Of importance for the analysis is only that its outcome is not representative for the error distribution otherwise seen in the data, and therefore SSDs represent half-fractions of full datasets having one outlier.

It has been stressed above that the only significant factor is factor 15. Williams [1] reported that in subsequent experiments he tried seven of the factors and found them all influential. That is of course not a contradiction, since *not significant* does not mean the same as *not influential*.

The outlier was more serious for SSDs than the negligense to transform the response variable. Since there is only one influential factor (15), its systematic effect can arbitrarily be characterized as additive or multiplicative, in combination with a response variance depending on the response level. A peculiar consequence is achieved, however, when in one of the studies in [7], extremely large artificial effects are added to the response. The consequence is that additivity of effects is prescribed, and that the log-transform is no longer even possible, since negative responses occur. Hence, the resulting structure is additive in factor effects but with a (residual) variance that depends (only) on the level of factor 15 — a fairly artificial situation.

Returning to the main topic, if the aim is not to show the influence of an outlier or to give only a mere illustration, the conclusion for use of Williams' data in SSD studies is:

- It is acceptable to use Lin's half, and other halves that exclude run #14, if also the underlying (true) structure of the full dataset is estimated with run #14 excluded (or replaced).

- It is not advisable to use Wang's or other halves that include run #14 [7–9], unless its response value is replaced by some more reasonable value.

References

1. Williams KR. Designed experiments. *Rubber Age* 1968; **100**: 65–71.

2. Box GEP, Draper NR. *Empirical Model Building and Response Surfaces.* John Wiley: New York, 1987.

3. Box GEP, Draper NR. *Response Surfaces, Mixtures, and Ridge Analysis.* John Wiley: New York, 2007.

4. Sudjianto A, Wasserman, GS, Sudarbo, H. Genetic subsets regression. *Computers Ind. Engng.* 1996; **30**: 839–849.

5. Lin DKJ. A new class of supersaturated designs. *Technometrics* 1993; **35**: 28–31.

6. Holcomb DR, Montgomery DC, Carlyle WM. Analysis of supersaturated designs. *J Quality Technology* 2003; **35**: 13–27.

7. Dejaegher B, Capron X, Vander Heiden Y. Generalized FEAR method to estimate factor effects in two-level supersaturated designs *J. Chemometr.* 2007; **21**: 303–323.

8. Abraham B, Chipman H, Vijayan K. Some risks in the construction and analysis of supersaturated designs. *Technometrics* 1999; **41**: 135–141.

9. Wang PC. Comments on Lin (1993). (With response) *Technometrics* 1995; **37**: 358–359.

10. Lu X, Wu X. A strategy of searching active factors in supersaturated screening experiments. *Journal of Quality Technology* 2004; **36**: 392–399.

11. Liu Y, Ruan S, Dean AM. Construction and analysis of Es^2 efficient supersaturated designs. *Journal of Statistical Planning and Inference* 2007; **137**: 1516–1529.

12. Westfall PH, Young SS, Lin DKJ. Forward selection error control in the analysis of supersaturated designs. *Statistica Sinica* 1998; **8**: 101–117.

13. Yamada S. Selection of active factors by stepwise regression in the data analysis of supersaturated designs. *Quality Engineering* 2004; **16**: 501–513.

14. Plackett RL, Burman JP. The design of optimum multifactorial experiments. *Biometrika* 1946; **33**: 305–325.

15. Joiner BL. Lurking variables: Some examples. *Amer. Statistician* 1981; **35**: 227–233.

16. Bisgaard S, Sutherland M. Split plot experiments: Taguchi's Ina Tile experiment reanalyzed. *Quality Engineering* 2003; **16**: 157–164.

17. Daniel C. *Applications of Statistics to Industrial Experimentation.* John Wiley: New York, 1976.

18. Box GEP, Meyer RD. Analysis of unreplicated factorials allowing for possibly faulty observations. In Mallows CL (editor): *Design, Data and Analysis.* Wiley 1987: 1–12.

19. Box GEP. Finding bad values in factorial designs. *Quality Engineeering* 1991; **3**: 405–410.

20. Sundberg R. Interpretation of unreplicated two-level factorial experiments, by examples. *Chemometrics Intell. Lab. Syst.* 1994; **24**: 1–17.

21. Carlson R. *Design and Opimization in Organic Synthesis, 2nd ed.* Elsevier, 2007.

Appendix: Table of data

Partially corrected, incomplete sets of data from Williams' experiment [1] are found in [2, 3] (exercise 5.20), [8], and [7] (Table II). A complete version is given here, with column 13 reconstructed and with the three interaction columns added. The latter are denoted A1–A3, where A3 is the defining contrast used by Lin [5].

Run

Run	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	A1	A2	A3	y
#1	+	+	+	−	−	−	+	+	+	+	+	+	+	−	+	+	+	+	+	+	−	+	+	+	−	+	−	133
#2	−	+	−	−	−	−	+	+	+	+	−	+	+	+	+	−	+	+	−	+	+	+	+	−	−	+	+	49
#3	+	−	−	−	−	−	+	+	+	−	−	−	+	+	+	+	−	+	−	−	+	+	+	−	+	−	−	62
#4	+	−	−	+	+	+	−	−	−	+	−	−	+	−	−	+	−	−	−	−	+	+	−	+	−	+	−	45
#5	+	+	−	−	+	+	−	−	−	−	+	−	−	+	−	+	−	+	−	+	−	−	−	−	+	−	+	88
#6	+	+	−	+	−	+	+	+	−	+	+	−	+	−	+	+	+	−	−	−	−	−	−	−	+	−	−	52
#7	−	−	+	+	+	+	+	−	+	+	−	+	−	+	+	−	+	+	+	+	+	+	−	−	+	+	+	300
#8	−	−	+	+	+	+	−	+	+	+	−	−	−	+	+	+	+	+	+	+	+	+	−	−	−	−	−	56
#9	−	−	+	+	+	+	−	−	−	−	+	+	+	+	−	−	+	+	+	+	+	−	−	−	+	+	−	47
#10	−	−	−	−	+	−	−	−	+	+	+	+	+	+	+	+	+	+	−	+	−	−	+	+	+	+	+	88
#11	+	−	+	+	+	+	+	−	+	−	+	+	−	−	+	−	+	−	+	−	−	+	+	−	−	−	−	116
#12	−	+	+	−	−	+	−	−	−	+	−	+	+	−	+	+	−	+	+	−	+	+	+	+	+	+	+	83
#13	−	+	+	+	+	−	+	+	+	+	−	+	+	−	+	−	+	−	+	−	+	+	+	+	+	+	−	193
#14	−	−	−	+	+	−	−	−	−	+	−	−	−	+	−	+	−	+	+	−	−	−	+	+	+	+	+	230
#15	+	−	+	+	+	−	+	+	+	+	+	+	+	+	+	+	+	−	+	+	−	−	+	+	+	+	+	51
#16	−	+	+	−	+	+	+	+	+	−	+	+	+	−	+	+	+	+	+	−	−	−	+	+	+	+	+	82
#17	−	−	−	+	+	+	−	−	−	−	+	+	−	−	+	+	−	+	+	−	+	+	−	−	−	−	−	32
#18	+	−	+	+	+	+	+	−	+	+	+	+	+	+	+	+	+	+	+	+	+	+	−	+	+	+	+	58
#19	+	−	+	+	+	+	−	+	+	−	+	+	−	+	−	−	−	−	−	−	−	+	−	−	+	+	+	201
#20	+	+	−	+	+	−	+	+	−	+	−	+	+	+	+	+	+	+	+	+	−	+	−	+	−	+	+	56
#21	−	+	−	−	−	+	−	−	+	−	−	+	−	+	−	−	−	−	+	−	−	+	−	+	+	−	+	97
#22	+	+	+	−	+	−	+	+	+	+	+	+	+	+	+	+	+	+	+	−	−	−	−	+	−	−	−	53
#23	−	−	−	−	−	−	+	+	+	−	+	−	−	−	−	−	−	−	+	−	+	+	+	+	−	−	−	276
#24	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	+	−	+	+	+	+	−	+	+	+	145
#25	+	+	−	−	−	−	+	+	+	−	+	−	−	+	−	−	+	−	−	−	+	+	+	−	−	+	+	130
#26	−	−	−	−	−	−	−	+	−	−	+	−	+	+	+	+	+	+	−	+	−	−	+	−	+	+	−	55
#27	+	−	−	+	−	−	+	+	−	+	+	−	−	+	−	+	−	−	−	+	−	+	+	−	+	+	+	160
#28	−	−	+	−	−	−	−	−	−	−	+	+	−	+	−	−	−	−	−	+	−	+	−	−	−	−	−	127